

Please replace the present Specification with the attached replacement Specification.

Please replace the present Abstract with the attached replacement Abstract.

- 5 **Please replace the present FIGs. 1, 4, and 5 with the attached replacement Figures 1, 4, and 5**

Copies of the replacement Figures with the changes marked in red have been provided; a letter to the Chief Draftsperson accompanies this response.

10

Remarks

The rejections of claims 112, 119, and 124 under 35 U.S.C. 112, second paragraph

The amendment to claim 112 has replaced the language "the request otherwise being executed in the first database system" with "the request being executed in the first database system when the query analyzer does not determine [that the request includes a specifier that cannot be interpreted in the first database system]" and has thereby overcome the rejection. Examiner will immediately see that claim 112 as amended is completely supported by the Specification as filed.

With regard to the rejection on the basis of a lack of antecedent for "second database" in claims 112, 119, and 124, Applicants' attorney respectfully points out that the claims use the language "a second database system of the plurality of database systems". There has been no earlier mention of any "second database system", so "a second database system" is perfectly proper; as for the "plurality of database systems", its antecedent may be found in the preambles of claims 112 and 119 and at line 4 of claim 124. There is thus no basis for rejecting these claims because of the lack of an antecedent.

The amendments to the Specification, Drawing, and Abstract

Examiner will immediately see that these amendments add no new matter to the Specification and Drawing, but instead serve merely to make the nomenclature used in the Specification consistent with that used in the Drawing. With regard to the page numbers of the Leverenz reference, they are correct in their present form. In the Leverenze reference, the pages in each chapter are independently numbered and the page

numbers consequently have the form <chapter number>-<page number>. The language objected to by Examiner, "pages 30-5 through 30-11" thus means "chapter 30, pages 5-11" and is correct.

5 The rejections of claims 112-1131 under 35 U.S.C. 103

Claims 112-131 are presently in the application. Examiner has rejected all claims as obvious over the combination of Draper and Jadav and Gupta. In the following, Applicants will show that Examiner has failed to make the *prima facie* case of obviousness required by MPEP 2142 for the following reasons:

- 10 1. because the Draper and Jadav and Gupta references do not, when combined, show all of the limitations of Applicants' claims.
2. even if the combined references are taken to show all of the limitations of Applicants' claims, it was not obvious at the time the invention was made to make such a combination.

15

Failure of the references to disclose all of the limitations of claim 112

MPEP 2142 requires that an examiner make a *prima facie* case of obviousness in order to reject a claim under 35 U.S.C. 103. A necessary part of the *prima facie* case is the citation of references which show every limitation of the claim under rejection. Claim

20 112 is exemplary for the independent claims in the above patent application. As presently amended, it reads as follows:

- 25 **112. (currently amended)** Apparatus for responding to a request, the request including one or more specifiers referring to objects belonging to a plurality thereof in a distributed database system that includes a plurality of database systems and the apparatus comprising:
- a first database system of the plurality of database systems;
- a query analyser that determines whether the request includes a specifier that cannot be interpreted in the first database system; and
- 30 a redirector which responds to the request when the query analyzer so determines by causing the request to be executed at least in part in a second database system of the plurality of database systems, the request being executed in the first database system when the query analyzer does not so determine.

35

The claim requires “a distributed database system” that includes “a first database system” and “a second database system.” What is meant by “data base system” is a system that works like a relational database system. In a relational database system, the data in the system is not located by means of an address (memory address, file name, URL, and the like), but by means of an SQL query which contains specifiers referring to objects in the database system. The database system executes the query and interprets the specifiers as it does so. So that the claim is not limited to relational database systems or queries in SQL, the term “request” has been used in the claim.

10 The claim further requires "a query analyser that determines whether the request includes a specifier that cannot be interpreted in the first database system" and “a redirector which responds to the request when the query analyzer so determines by causing the request to be executed at least in part in a second database system of the plurality of database systems”. If the request’s specifiers can be interpreted in first database system, the request is executed there.

Neither reference discloses the claim’s “distributed database system”. As pointed out above, as used in the claim, the term “database system” means a system that references data by means of requests containing specifiers for objects in the database system rather than addresses and locates the data by executing the request and interpreting the specifiers contained in the request. Though Draper uses the term “distributed database”, there is no disclosure whatever in the reference that indicates that any of the components of his “distributed database” references data by means of requests that contain specifiers or that the components of the “distributed database” interpret the specifier as they execute the request. Indeed, Draper’s only example, disclosed at col. 9, lines 33-53, discloses caches which contain HTML versions of documents.

As for Jadav and Gupta, FIGs. 3, 4, 6, and 7 do not show a “distributed database system”. Instead, they show a single database system. FIG. 4 shows a Web driver that builds SQL statements in addition to the database system; FIG. 6 shows a Web server with a cache for large objects; FIG. 7 shows operation of the cache of FIG. 6. The large objects in the

cache are referred to by what Jadav and Gupta call “static queries”; they are defined at page 15, col. 1, lines 36-38:

A *static* query for a large object has the large object handle hardwired into the query

5

Thus, in the terminology used in the present discussion, Jadev and Gupta’s static query is an address. A large object may also be referred to by a *dynamic* query, in which “the value of the large object handle is obtained at run-time” (page 15, col. 1, lines 38-39), (i.e., by executing the query). A dynamic query is thus an example of a “request” as that term is used in claim 1. As set forth at page 17, col. 1, lines 29-41, the cache determines whether the query is static or dynamic; if it is dynamic, the cache always passes it on to the database system; if it is static, the cache looks for the large object in the cache; if it is not there, it passes the static query on to the database system. Jadav and Gupta’s cache consequently clearly cannot be understood to be a “database system” as that term is defined in Applicants’ claim 1.

15

Because neither reference shows Applicants’ distributed database system, with its “first database system” and “second database system, the references when combined do not show the claimed “first database system”, “second database system”, “query analyzer”, and “redirector which responds to the request when the request includes a specifier that cannot be interpreted in the first database system by causing the request to be executed at least in part in a second database system of the plurality of database systems, the request otherwise being executed in the first database system”. In particular, because there is nothing corresponding to Applicants’ “first database system” in either reference, Jadev and Gupta’s redirection mechanism cannot “respond to the request when the request includes a specifier that cannot be interpreted in the first database system by causing the request to be executed at least in part in a second database system”. Further, because there is nothing corresponding to the “first database system” in either reference, the combination of references cannot show the limitation, “the request being executed in the first database system when the query analyzer does not so determine”.

20

25

30

Because the references, when combined, do not disclose all of the limitations of Applicants' claim 1, Examiner has not made her *prima facie* case for the rejection of the claim. The arguments set forth above with regard to claim 112 can be applied *mutatis mutandis* to claims 119, 124, 125, 128, and 131 as well.

5

Traversal of Examiner's assertion that redirecting a query from a first database system to a second database system is obvious

Even if the combined references are taken to disclose all of the limitations of Applicants' invention, Examiner must show that it would have been obvious at the time the invention was made to combine them. Examiner's showing that it would have been obvious to one of ordinary skill in the art to combine Draper and Jadav and Gupta reads as follows: "[the modification of Draper by adding Jadav and Gupta's redirector is obvious] because such a modification would allow Draper to serve stored documents locally (rather than in the database) ..." (Office action of 2/6/2006, page 9, lines 6-13)

15

This amounts to the advantage offered by any kind of caching. The problem with Examiner's argument is that *how* to make a cache where the cache was a database system and consequently could execute requests by interpreting specifiers in the requests was not obvious at the time the invention was made. The history of computing provides ample proof that it was not obvious. Relational database systems go back to the early '70's, and are thus over thirty years old. Caches go back at least as far, beginning with memory caches, continuing with caches for files, and today including caches for Web pages.

During the last thirty years, many caches have been based on the principles of *transparency* and *redirection*. A cache is *transparent* if a program referring to data which may be in the cache refers to the data in the same way whether it is in the cache or elsewhere. See page 6, line 11 of the application as filed. A cache performs *redirection* if it automatically transfers a reference to data that is not contained in the cache to a source of the data. Thus, a PC browser maintains a cache of HTML pages in the PC. When a user specifies an HTML page by URL, the browser first looks in the cache for the page; if it is not there, the browser redirects the URL to the Internet.

As pointed out above, both caches and relational database systems have existed for over thirty years. Further, as Examiner points out in her rejection, caches provide useful advantages. Nevertheless, in the five years and 11 months since the mailing of the first
5 Office action in this application on 6/20/00, neither of the two Examiners who have worked on the application have found a reference which shows a database system that could execute requests by interpreting specifiers of objects in the database system, which was transparent to the application program making the request, and which redirected the request to another database system if the specifier could not be interpreted. The first
10 database system of Applicants' claim 112 is such a database system, and as such, it can serve as a cache.

Given the thirty-year coexistence of caches and relational database systems, the lack of any such reference shows beyond a doubt that constructing a system like the one set forth
15 in claim 112 was not obvious to those who were acquainted with caches, to those who were acquainted with database systems, or to the many people in the computer arts who were acquainted with both.

To understand why Applicants' invention of claim 112 was not obvious to those skilled
20 in the relevant arts, one need only look at the difference between the way data is referenced in a database system and the way data in memory, files in a file system, or HTML pages in the Internet are referenced. Locations in memory, files in a file system, and HTML pages in the Internet are all referenced by means of *addresses*, i.e., there is a value which, in the context in which the addressing is being done, uniquely identifies the
25 memory location, the file, or the HTML page. In the case of locations in memory, it is a memory address; in the case of a file, it is the pathname for the file; in the case of the Web page, it is the page's URL. The address of the item is relatively persistent: the address of an item in memory may be valid for as long as the execution of the process lasts to which the memory belongs; the address of an item in the file system may last as
30 long as the file system exists; the address of an item in the Internet may last as long as the Internet exists. Because the address of an item is relatively persistent, when the item is

placed in a cache, the address of the item can be mapped onto the location of the item in the cache. The mapping is typically done using techniques such as hash tables: the address of the object is hashed to obtain the index of an element in the hash table and that element or an element linked to that element contains the location of the item in the
5 cache.

The requests which are used to locate data in the system of claim 12 are not addresses. There is no persistent relationship between the text of a request and the data it returns. Instead, the results of the execution of a request may vary from execution to execution of
10 the request and thus cannot be determined until the request is actually executed. For example, if a query's WHERE clause specifies a condition, the results of the execution will depend upon the condition. Because there is no persistent relationship between the text of a query and what it returns, a mapping cannot be established between the text of the request and a location in a cache that is a database system. Instead, completely
15 different techniques must be used. The techniques that are used to do this in the embodiment of the invention of claim 12 which is disclosed in the above application are disclosed beginning at page 11, line 31.

The only conclusion that can be drawn from a comparison of the discussion beginning at
20 page 11, line 31 with the techniques generally used in caches to map addresses onto cache locations is that it is by no means obvious from these techniques for mapping addresses onto cache locations how to deal with requests in a database system that is a cache. This conclusion is reinforced by the Jadav and Gupta reference, which states at page 19, col. 1, lines 1-12:

In this paper we did not address the issue of caching large objects when they are accessed using dynamic queries. The reason for this is as follows : dynamic queries are always embedded in a HTML page, within MYSQL tags. Webdriver/Webcache receives the client query for a page or a large object, hence caching can be done for static queries. Dynamic queries are always parsed, processed and formatted by the WebExplode function, which resides on the database server. If caching were to be implemented for dynamic queries, network traffic would be added between the database and the web server after the query has been processed.

The reason that Jadav and Gupta believe extra network traffic would be required is that they cannot conceive of a cache that would be capable of interpreting dynamic queries, and consequently believe that all dynamic queries must be sent to the source database system for interpretation. The reference thus not only does not support the combination used by Examiner to reject the claims, but actually teaches against the combination.

Conclusion

Applicants have amended claims 112 and 125 to overcome the rejections of those claims under 35 U.S.C. 112, 2. par., have traversed the remaining rejections under 35 U.S.C. 112, have amended their Specification, Drawing, and Abstract to overcome the objections thereto, and have shown that the combination of Draper and Jadav and Gupta does not disclose all of the limitations of Applicants' claim 112; Applicants have further shown that even if the combination is taken to disclose all of the limitations of Applicants' claim 112, it was by no means obvious to those skilled in the computer arts to make the combination at the time the invention was made. Applicants' invention of claim 112 is thus patentable over the references either because the combination of Draper and Jadav and Gupta does not disclose all of the limitations of Applicants' claim 112 or because the combination is not obvious. Examiner will immediately see that these arguments concerning claim 112 apply equally to Applicants' other independent claims.

Applicants have thus been fully responsive to Examiner's Office action of 2/6/02, as required by 37 C.F.R. 1.111(b) and respectfully request that Examiner continue with her

examination and allow the claims as amended, as provided in 37 C.F.R. 1.111(a). No fees are believed to be required by way of this amendment. Should any be, please charge them to deposit account number 501315.

5

Respectfully submitted,


Attorney of record,

10

Gordon E. Nelson
57 Central St., P.O. Box 782
Rowley, MA, 01969,
Registration number 30,093
Voice: (978) 948-7632

15

Fax: (866) 723-0359

May 5, 2006
Date**Certificate of Mailing**

20 I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

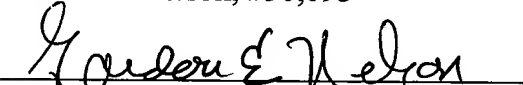
25

on May 5, 2006

(Date)

Gordon E. Nelson, #30,093

30


(Signature)

MARKED UP COPY

Web servers with queryable dynamic caches

Background of the invention

5

1. Field of the invention

The invention concerns caching of data in networks generally and more specifically concerns the caching of queryable data in network servers.

10 2. Description of the prior art

Once computers were coupled to communications networks, remote access to data became far cheaper and easier than ever before. Remote access remained the domain of specialists, however, since the available user interfaces for remote access were hard to learn and hard to use. The advent of World Wide Web protocols on the Internet ~~have~~has finally made remote access to data available to everyone. A high school student sitting at home can now obtain information about Karlsruhe, Germany from that city's Web site and a lawyer sitting in his or her office can use a computer manufacturer's Web site to determine what features his or her new PC ought to have and then configure, order, and pay for the PC.

20

A consequence of the new ease of remote access and the new possibilities it offers for information services and commerce has been an enormous increase in the amount of remote access. This has in turn lead to enormous new burdens on the services that provide remote access and the resulting performance problems are part of the reason why the World Wide Web has become the World Wide Wait.

25

FIG. 1 shows one of the causes of the performance problems. At 101, there is shown the components of the system which make it possible for a user at his or her PC to access an information source via the World Wide Web. Web browser 103 is a PC which is running Web browser software. The Web browser software outputs a universal resource locator (URL) 104 which specifies the location of a page of information in HTML format in the

30

World Wide Web and displays HTML pages to the user. The URL may have associated with it a message containing data to be processed at the site of the URL as part of the process of obtaining the HTML page. For example, if the information is contained in a database, the message may specify a query on the ~~data-based~~database. The results of the query would then be returned as part of the HTML page. Internet 105 routes the URL 104 and its associated message to the location specified by the URL, namely Web server 107. There, HTML program 109 in Web server 107 makes the HTML page 106 specified by the URL and returns it to Web browser 103. If the message specifies a query on the database in database server 115, HTML program 109 hands the message off to Web application program 111, which translates the message into a query in the form required by data access layer 112.

Data access layer 112 is generally provided by the manufacturer of database server 115. It takes queries written in standard forms such as OLE-DB, ODBC, or JOBC, converts the queries into the form required by database server 115, and places the queries in messages in the form required by network 113. Database server 115 then executes the query and returns the result via network 113 to data access layer 112, which puts the results into the required standard form and returns them to ~~Web-application~~Web application program (WEB APP) 111, which in turn puts the result into the proper format for HTML program 109. HTML program 109 then uses the result in making the HTML page 106 to be returned to browser 103.

As may be seen from the above description, a response to a URL specifying a page whose construction involves database server 115 requires four network hops: one on Internet 105 from browser 103 to Web server 107, one on network 113 from server 107 to server 115, one on network 113 from server 115 to server 107, and one on Internet 105 from server 107 to browser 103. If more than one query is required for an HTML page, there will be a round trip on network 113 for each query.

Moreover, as shown at 117, a typical Web transaction is a series of such responses: the first HTML page includes the URL for a next HTML page, and so forth. The

transaction shown at 117 begins with a request for an HTML page that is a form which the user will fill out to make the query; ~~data-based~~database server 115 provides the information for the HTML page. When that page is returned, the user fills out the form and when he or she is finished, the browser returns a URL with the query from the form to server 107, which then deals with the query as described above and returns the result in another HTML page. That page permits the user to order, and when the user orders, the result is another query to database server 115, this time, one which updates the records involved in the transaction.

Not only do Web transactions made as shown in FIG. 1 involve many network hops, they also place a tremendous burden on ~~data-based~~database server 115. For example, if ~~data based~~database server 115 belongs to a merchant who sells goods on the Web and the merchant is having a special, many of the transactions will require exactly the same sequence of HTML pages and will execute exactly the same queries, but because system 101 deals with each request from a web browser individually, each query must be individually executed by database server 115.

The problems of system 101 are not new to the designers of computer systems. There are many situations in a computer system where a component of the system needs faster access to data from a given source, and when these situations occur, the performance of the system can be improved if copies of data that is frequently used by the component are kept at a location in the system to which the component has faster access than it has to the source of the data. When such copies exist, the location at which the copies are kept is termed a *cache* and the data is said to be *cached* in the system.

Caching is used at many levels in system 101. For example, browser 103 keeps a cache of previously-displayed HTML pages, so that it can provide a previously-displayed HTML page to the user without making a request for the page across Internet 105. Web server 107 similarly may keep a cache of frequently-requested HTML pages, so that it can simply return the page to the user, instead of constructing it. Database server 115, finally, may keep a cache of the information needed to answer frequently-made queries,

so that it can return a result more quickly than if it were starting from scratch. In system 101, the most effective use of caching is in Web server 107, since data that is cached there is still accessible to all users of internet 105, while the overhead of the hops on ~~data~~ accessnetwork 113 is avoided.

5

Any system which includes caches must deal with two problems: maintaining consistency between the data in the cache and the data in the data source and choosing which data to cache. In system 101, the first problem is solved in the simplest way possible: it is the responsibility of the component using the data to determine when it needs a new copy of the data from the data source. Thus, in browser 103, the user will see a cached copy of a previously-viewed HTML page unless the user specifically clicks on his browser's "reload" button. Similarly, it is up to HTML program 109 to determine when it needs to redo the query that provided the results kept in a cached HTML page. The second problem is also simply solved: when a new page is viewed or provided, it replaces the least recently-used cached page.

10

15

Database systems such as the Oracle8™ server, manufactured by Oracle Corporation and described in Leverenz, et al., *Oracle8 Server Concepts*, release 8.0, Oracle Corporation, Redwood City, CA, 1998., move a copy of a database closer to its users by *replicating* the original database at a location closer to the user. The replicated ~~data-base~~ database may replicate the entire original or only a part of it. Partial replications of a database are termed *table snapshots*. Such table snapshots are read-only. The user of the partial replication determines what part of the original database is in the table snapshot. Consistency with the original database is maintained by *snapshot refreshes* that are made at times that are determined by the user of the table snapshot. In a snapshot refresh, the table snapshot is updated to reflect a more recent state of the portion of the original database contained in the snapshot. For details, see pages 30-5 through 30-11 of the Leverenz reference.

20

25

There are many applications for which the solution of letting the component that is doing the caching decide when it needs a new page causes problems. For example, when the

30

information in a data source is important or is changing rapidly (for example, stock prices), good service to the user requires that the information in the caches closely tracks the information in the data source. Similarly, there are many situations where caching all data that has been requested causes problems. For instance, in a cache run according
 5 to least recently-used principles, any HTML page that is produced by HTML program 109 or received in browser 103 is cached and once cached, stays in the cache and takes up space that could be used for other HTML pages until it attains least recently-used status.

10 | When Web server 107 includes a ~~Web application~~Web application program 111 involving a database server 115, there is still another problem with caching in web server 107: since the data is cached in the form of HTML pages, it is not in *queryable* form, that is, a cached HTML page may contain data from which another query received from Web browser 103 could be answered, but because the data is contained in an HTML page
 15 instead of a database table, it is not in a form to which a query can be applied. Thus, even though the data is in server 107, server 107 must make the query, with the accompanying burden on ~~data-based~~database server 115 and delays across network 113, and the HTML page containing the result of the query must be separately cached in server 107.

20

What is needed to solve these problems is a web server 107 that has a cache in which cached data is to the extent possible in queryable form, in which the cached data is dependably updated when the data in the source changes, and in which selection of data from a source for caching is based on something other than the mere fact that a URL
 25 received from a web browser referenced the data. It is an object of the invention disclosed herein to provide servers and data sources that solve the above problems.

Summary of the invention

The problem of updating the server's cache is solved by having the sources of the cached
 30 information send update messages to the server each time the cached information changes in the information source. The problem of determining what to cache is solved by

determining what to cache on the basis of probable future requests for the information. The determination of what information will probably be made the subject of future requests can be made in the server, in the data source, or elsewhere.

- 5 The problem of queryable data is solved by using a database system in the server as the cache. If the information necessary to run the query is present in the cache database system, the query is run on the cache database system; otherwise, it is run on a source database system. The cache database system is made transparent to application programs running on the server by setting up the data access layer so that it can run queries on
10 either the source database system or the cache database system. The data access layer receives a query in standard form from the application program; it then determines whether the information needed for the query is present in the cache database; if it is, the data access layer runs the query on the cache database; if it is not, the data access layer runs the query on the source database system.

15

- In a further aspect of the invention, the standard form of the query uses global dataset identifiers, while the copies of the datasets in the cache database use local dataset identifiers. A query analyzer in the cache database receives the global dataset identifiers used in the query from the data access layer; if copies are present in the cache, the query
20 analyzer indicates that to the data access layer and returns the local dataset identifiers for the copies to the data access layer. The data access layer then uses the local dataset identifiers to query the cache database.

- Other objects and advantages will be apparent to those skilled in the arts to which the
25 invention pertains upon perusal of the following *Detailed Description* and drawing, wherein:

Brief description of the drawing

- FIG. 1** is an example of a prior-art system for performing queries via the World Wide
30 Web;

FIG. 2 is a high-level block diagram of a system of the invention;

- FIG. 3 is a detailed block diagram of details of an implementation of a server 203(i);
 FIG. 4 is a detailed block diagram of details of an implementation of source database server 237;
 FIG. 5 is a detail of cache database description 305; and
 5 FIG. 6 is a flowchart of the operation of query dispatcher 351

Reference numbers in the drawing have three or more digits: the two right-hand digits are reference numbers in the drawing indicated by the remaining digits. Thus, an item with the reference number 203 first appears as item 203 in FIG. 2.

10

Detailed Description

The following *Detailed Description* will begin with a conceptual overview of the invention and will then describe a presently-preferred embodiment of the invention.

15 Overview of the invention: FIG. 2

FIG. 2 shows a system 201 for retrieving information via a network which includes one or more network servers 203(0..n). A server 203(i) and another server 203(n) are shown in FIG. 2. Each server 203(i) includes a queryable cache 219 that is automatically
 20 updated when ~~information cached in cache~~cached data 223 changes in source database 241 and in which the contents of cached data 223 are determined by an analysis of what queries will most probably be made by users of server 203(i) in the immediate future. ~~Each of Server servers 203(i..n) is a Web server 107 as shown in FIG.1, and thus has an HTML component 109, a Web application component programs 111, and a data access component layer 253 which is a version of data access component layer 112 which has~~
 25 been modified to work with queryable cache 219. Server 203 could, however, communicate with its users by any other kind of network protocol. Server 203 further communicates with source ~~data based~~database (DB) server 237 by means of network 113, which may use any protocol which is suited to the purpose.

FIG. 2 shows one of servers 203(0..n), server 203(i), in detail. As before, ~~Web application~~Web application program 111 provides a query in a standard form to ~~data access~~data access layer 253. Here, however, ~~data access~~data access layer 253 has access not only to source database server 237 via network 113, but also to queryable cache 219, which contains a cache ~~data-based~~database 236 ~~that has a~~whose cached data 223 is a copy 223 of a portion of the data in source database 241. When ~~data access~~data access layer 253 receives a query from ~~web application~~Web application program 111, it first presents the query to queryable cache 219, as shown at Q 215. If cached data 223 includes the data specified in the query, queryable cache 219 returns result (R) 217, which ~~data access~~data access layer 253 returns to ~~Web application~~Web application program 111. If cached data 223 does not include the data specified in the query, queryable cache 219 returns a hit or miss (H/M)~~miss signal (M)~~ 216 indicating a miss to ~~data access~~data access layer 253, which then makes the query via network 113 to source database server 237 and when it receives the result, returns it to ~~Web application~~Web application program 111. The query made in response to the miss signal appears as miss query (MQ) 224 and the response appears as miss response (MR) 226.

It is important to note here that because the interactions with queryable cache 219 and with source database server 237 are both performed by data access layer 253, the existence of queryable cache 219 is completely *transparent* to ~~Web application~~Web application program 111. That is, a ~~Web application~~Web application -program 111 that runs on Web server 107 will run without changes on Web server 203(i).

Continuing in more detail with queryable cache 219, the data cached in queryable cache 219 is contained in cache database 236, which, like any database, contains data, in this case, copies of datasets (database tables) from source database 241 that are cached in queryable cache 219, and a query engine (QE) 221, which runs queries on the datasets in cached data 223. The portion of queryable cache 219 which receives queries from data access layer 253 is data access (DA) interface 212. Data access interface 212 has two functions:

- It determines whether the query can be executed on cached data 223 required to execute query 215 and generates ~~miss signal~~H/M signal 216 indicating a miss if it does not.
- If cached data 223 does contain the data, it puts query 215 into the proper form for cache database 236.

Data access interface 212 makes the determination whether the query can be executed by analyzing the query to determine the query's *context*, that is, what datasets are required to execute the query and then consulting a description of cached data 223 to determine whether these datasets are present in cached data 223. The datasets are specified in the query by means of dataset identifiers, and consequently, the context is for practical purposes a list of the identifiers for the required data sets. The description 223 of course includes the dataset identifiers for the cached data sets. If the required datasets are present, data access interface 212 makes cache query (CQ) 245, which has the form required to access the data in cache ~~data-based~~database 236. Cache database 236 returns cache result (CR) 247, which data access interface 212 puts into the form required for result (R) 217.

Because cached data 223 is contained in cache database 236, cached data 223 is *queryable*, that is, if a dataset is contained in cached data 223, queryable cache 219 can return as a result not only the entire dataset, but any subset of that dataset that can be described by a query. For example, if cached data 223 includes a dataset that lists all of the kinds of shirts sold by a company engaged in Web commerce and the list of kinds includes the colors that each kind of shirt is available in, queryable cache 219 will be able to handle a query for which the result is a list of the kinds of shirt that are available in red.

Cached data 223 is kept consistent with source database 241 by means of update transmitter (UPDATE XMIT) 243 in source database server 237 and update receiver 210 in queryable cache 219. Whenever a change occurs in source database 241 in a dataset of which there may be a copy in cached data 223, update transmitter 243 generates a cache update query (CUDQ) 234 specifying the change and sends CUDQ 234 via network 113 to each of servers 203(0..n). Update receiver 210 receives CUDQ 234 from network 113

and determines from the data set description maintained by ~~DA~~data access interface 212 whether the dataset is in face in cached data 223; if it is, it puts the cache update query into the proper form for cache database 236 and provides it to cache refresher 249, which then runs cache update query (CUDQ) 251 on cache database 236.

5

Data set manager (DSM) 213 decides generally what copies of datasets from source database server 237 are to be included in cache database 236. The information that DSM 213 uses to make this determination is contained in query information (QI) 208. Query information 208 may be any information available to server 203(i) which can be used to predict what datasets of source database 241 will most probably be queried in the near future. For example, if a company engaged in Web commerce is having a 1-day sale on certain items for which there are datasets in source database 241, query information 208 may indicate the datasets for the items and the time of the 1-day sale. Using that information, DSM 213 can obtain the datasets from source database 241 and cache them in cache database 236 before the beginning of the sale and remove them from cache database 236 after the end of the sale.

10

15

Another kind of query information 208 is a query log, a time-stamped log of the queries received from data access layer 253; if the log shows a sharp increase in the occurrence of queries for a given dataset, DSM 213 should cache the datasets for that query in queryable cache 219 if they are not there already. Conversely, if the log shows a sharp decrease in the occurrence of such queries, DSM 213 should consider removing these datasets from queryable cache 219. When DSM 213 determines that a dataset should be added to queryable cache 219, it sends a new data query (NDQ) ~~218-220~~ via network 113 to source ~~data-based~~database 241 to obtain the new data and when DSM 213 has the response (NDR ~~220~~218), it sends a delete query to query engine - 221 indicating the data to be deleted in cached data 223 to make way for the new data and then sends a cache update query 251 to cache refresher 249 to update the cache.

20

25

30

Data set manager 213 and query information 208 may also be implemented in part in source ~~data-based~~database server 237 or anywhere where information about the probability

of future queries may be obtained. When implemented in source ~~data-based~~database server 237, the query log would log each query 234 to source database 241 and at least the portion of data set manager 213 which reads the query log to determine what new data needs to be cached would be in source database server 237; when it determined that new data needed to be cached, it would send an update query with the new data to each of the servers 203. The component of DSM 213 that determines what is to be removed could also be in source database server 237, in which case, all queryable caches 219 would contain the same data in cached data 223, or that component could be in each server 203(i), with the component making decisions concerning what data to remove to accommodate the new ~~data-based~~database on the present situation in server 203(i). In such an arrangement, there can be a local query log in each server 203 in addition to the global query log in source database server 241~~237~~. Such an arrangement would permit different servers 203 to have different-sized queryable caches ~~223~~219; it would also permit different servers 203 to take local variations in the queries they are receiving into account in determining what data to remove from queryable cache 219. One way such variations might occur is if system 201 were set up so that different servers 203 preferentially received queries from users in different geographical locations.

FIG. 2 shows only a single source database server 237; there may of course be more than one; moreover, source database server 237 need not be a classical database system. Server 203(i) can be set up to be used with data sources containing any kind of queryable data, where *queryable* is defined as having a form which can be represented as a set of numbered rows of data. Such a set of numbered rows is termed a *rowset*. Database tables are of course one example of rowsets; others are files of data records, text files, and still and moving image data. If server 203(i) is used with data sources having only a single kind of queryable data, queryable cache 219 need only be set up to deal with that kind of queryable data.

If server 203(i) is used with data sources having more than one kind of queryable data, cache database 236 may be set up using a rowset representation that will accommodate all of the different kinds of queryable data. In that case, ~~DA~~data access interface 212,

DSM 213, and update receiver 210 will translate between the results and update queries received from the various data sources and the representations used in cache database 236. In other embodiments, there may be more than one cache database 236 in queryable cache 219, with different cache databases being used for different kinds of queryable data. Again, ~~DA~~data access interface 212, DSM 213, and update receiver 210 will perform the necessary translations.

Details of a preferred embodiment of a data access ~~layer~~ component ~~layer~~ and a queryable cache: FIGs. 3, 5, and 6

FIG. 3 shows a preferred embodiment 301 of data access layer 349 and queryable cache 302. Corresponding components of FIGs. 2 and 3 have the same names. Cache database 347 in embodiment 301 is an Oracle8 Server, which is described in detail in Leverenz, et al., *Oracle8 Server Concepts*, release 8.0, Oracle Corporation, Redwood City, CA, 1998. In preferred embodiment 301, ~~Web application~~Web application program 111 uses global data set identifiers in queries. The ~~Web application~~Web application programs 111 in all of the servers 203 use the same set of global data set identifiers. A cache ~~data based~~database 347 in a given server 203 has its own set of local data set identifiers for the data sets cached in cache ~~data based~~database 347. In preferred embodiment 301, then, one may speak of *global* queries and query contexts that use global data set identifiers and *local* queries and query contexts that use local data set identifiers. In the preferred embodiment, query analyzer 313 uses cached ~~data based~~database description (CDB DESC) 305 to translate global query contexts into local query contexts.

Data access layer 349 includes a new component, query dispatcher 351, which is the interface between data access layer 349 and queryable cache 302. FIG. 6 is a flowchart 601 of the operation of query dispatcher 351 in a preferred embodiment. Reference numbers in parentheses refer to elements of the flowchart. When data access layer 349 is preparing to query source database 241, it provides the global context for the query to query dispatcher 351 (605), which in turn provides global context 318 (FIG. 3) to query analyzer 313 (607). Query analyzer 313 determines whether the datasets identified by the global context are cached in cache database 347; if they are not, query analyzer 313

reports a miss 319 to query dispatcher 351 (609), which indicates to data access layer 349 that it is to place the global query (GQ) 353 on network 113.

If the datasets identified by the global context are cached in cache database 347, query analyzer 313 indicates that fact to query dispatcher 351 and also provides query dispatcher 351 with local context 316 for the datasets in cache database 347 (615). Query dispatcher 351 then provides the local context to data access layer 349, which uses the local context to make a local query 317 corresponding to the global query and then uses the local query to obtain local result 320 from cache database 347. It should be noted here that the operations involved in the translation from the global query to the local query and applying the local query to cache database 347 may be divided among data access layer 349, query dispatcher 351, and query analyzer 313 in many different ways; the advantage of the technique of flowchart 601 is that data access layer 349 can employ the same mechanisms to make local queries as it does to make global queries. All query analyzer 313 and query dispatcher 351 need do is supply data access layer 349 with the local context needed to make the local query.

Continuing with the details of queryable cache 302 and beginning with DA interface 304, interface 304 receives a global context 318 from query dispatcher 351 and depending on whether the datasets for the queries are in cache database 347, provides either local context 316 or a miss signal 319. DA interface 304 has two main components: query analyzer 313 and cache database description manager 303.

Query analyzer 313 analyzes global contexts received from data access layer 253 and other components of embodiment 301 to obtain the global context's global dataset identifiers. Having obtained the global dataset identifiers, query analyzer 313 provides them to CDB description manager 303, which looks them up in cache database description 305. Cache database description 305 is a table of datasets. At a minimum, there is an entry in the table for each dataset that has a copy in cache database 347. Each such entry contains the dataset's global identifier and its local identifier. The table also contains query information (QI) 307. CDB description manager 303 then returns an

indication of whether the dataset is in cache database 347 (H/M 311). If it is not, the query cannot be run on cache database 347, but must be run on source database 241, and consequently, query analyzer 313 returns a miss signal 319 to query dispatcher 351. If the query can be run on cache database 347, query analyzer 313 returns a hit signal 319 and also returns local context 316 for the query. As indicated above, query dispatcher 351 then provides local context 316 to data access layer 349, which uses it to make local query 317 on cache database 347. Cache database 347 then returns local result 320 to data access layer 349.

FIG. 5 shows details of CDB description 305. In a preferred embodiment, it is a table which has at least an entry 501 for each dataset of source database 241 of which here is a copy in cache database 347. Each entry 501 contains the global dataset identifier for the data set, by which the dataset is known in all servers 107 with queryable caches 219 containing copies of the dataset, the local data set identifier 505, by which the dataset is known in cache database 347, and number of queries information 507, which indicates the number of times the dataset has been queried over an interval of time. In the preferred embodiment, number of queries information 507 embodies query information 307.

An entry 501(i) for a given dataset is accessed in a preferred embodiment by a hash function 503513, which takes global dataset ID (GLOBAL DSID) ~~507~~511 for the dataset and hashes it into an entry index 509 in table 305. CDB description manager 303 then searches table 305 for the entry 501 whose field 503 specifies global DSID 511 beginning at entry index 509. If no such entry is found, the dataset is not in cache database 347 and CDB description manager 303 signals a miss 311 to query analyzer 313. Table 305 may also include entries 501 for global datasets that are not presently cached in cache database 347; in such entries, local dataset ID 505 has a null value and a miss is returned in response to the null value. The purpose of such entries is to maintain number of queries information 507 for such data sets, so that dataset manager 323 can determine whether to add the entry's dataset to cache database 347.

Update Rcvr 321 receives update queries provided by source database server 237 from data access 253 and uses query analyzer 313 to determine whether the dataset affected by the update is in cache database 347. If it is not, update rcvr 321 ignores the update; otherwise, it places local update query 329 in change queue 333. Refresher 331 reads
 5 queue 333 and executes its queries.

Data store manager 323 uses query information 307 in CDB description 305 to determine what datasets to add to or delete from cache database 347. With datasets to be added, DSM 323 makes the necessary queries to source database 241 and when the results
 10 arrive, DSM 323 makes them into update queries 329 and provides the update queries 329 to change queue 333, from which they are executed by refresher 331 as described above. DSM 323 further updates CDB description 305 as required by the changes it makes in cache database 347, as shown at 327.

15 In a preferred embodiment, DSM 323 and refresher 331 have their own threads or processes. It should also be pointed out here that CDB description 305 and change queue 333 could be implemented as database tables in cache database 347. Because these components are implemented independently of cache database 347 and because DA Interface 304 is used as an interface to cache database 347, embodiment 301 is to a large
 20 extent independent of the particular kind of database system employed to implement cache database 347. In embodiment 301, data access layer 349 only provides read queries to data access interface 304. All update queries go directly to server 237, without the update being entered in cache database 347. In other embodiments, queryable cache 219 may be implemented as a *writethrough* cache, i.e., the update may be entered in
 25 cache database 347 and also sent to server 237. It should be pointed out here that most ~~Web application~~Web application programs are *mostly-read* applications, that is, a Web user typically spends far more time reading information than he or she does changing it. For instance, in Web commerce, the “shopping” is mostly a matter of reading HTML pages, with updates happening only when the user adds something to his or her
 30 “shopping cart” or makes his or her purchases. In a system such as system 201, only making the purchases would typically involve an update of source database 241.

Details of source database server 237: FIG. 4

FIG. 4 shows a preferred embodiment of source database server 237. Source database server 237 in the preferred embodiment is implemented by means of an Oracle8 server
 5 | executing on a computer system that includes a source server disk drive 421 on which is
 stored source database 241 and source server memory 415 which contains buffer cache
 407 for copies of data values 421-422 from source database 241 and dictionary cache 409
 for copies of *metadata* from source database 241. Metadata is database tables whose
 contents describe the data in the database. Writebacks of cached data in source server
 10 | memory 415 to source database 241 are handled by database write process 325. Each of
 processes 401(0..n) represents and corresponds to a server 203(0..n) and handles queries
 resulting from cache misses, update queries, and queries from DSM 323 in the
 corresponding server 203. Two such processes, 401(i) and 401(n), are shown in FIG. 4.
 Dispatcher 311 gives each of these processes in turn access to shared server process 317,
 15 | which performs the actual queries and returns the results to the querying process, which
 in turn returns the results via network 235 to its corresponding server 203.

The Oracle8 implementation of source database server 237 is a standard Oracle8 database
 system to which has been added an implementation of update transmitter 243 which
 20 | automatically sends an update to queryable cache 219 in each of the servers 203(0..n)
 when data in source database 241 that has been copied to cached data 223 changes. The
 components of updater 243 in FIG. 4 are labeled with the reference number 243 in
 addition to their own reference numbers. The implementation of updater 243 in the
 preferred embodiment employs database *triggers*. A database trigger is a specification of
 25 | an action to be taken if a predefined change occurs in a data value or an item of metadata
 in the database. Definitions of the triggers are kept on source server disk 421 in trigger
defs 419; when a trigger is in use, it is in source server memory 415, as shown at trigger
defs 411. Many database systems permit definition of triggers; triggers in the Oracle8
 database system are described in detail at pages 17-1 through 17-17 of the Leverenz
 30 | reference.

In the preferred embodiment, when a process 401(i) corresponding to a server 203(i) receives a query from DSM 323 in server 203(i) for data to be added to server 203(i)'s cached data 223, process 203(i) executes set trigger code 403. This code sets an Oracle8 AFTER row trigger in metadata 417 for each row of data and/or metadata specified in the query. Shared server process 317 takes the action specified in the trigger whenever the trigger's row of data has been modified. The action specified for the trigger is to send a message to each of the servers 203(0..n) with an update query that modifies the data in cached data 223 in the same fashion as it was modified in source database 241. In the preferred embodiment, the action performed by the trigger is to place the message with the update query in message queue 414, which is implemented as an Oracle8 advanced queue. Message queue 414 is read by update process 402, which sends the messages in queue 414 to each of the servers 203(0..n).

Adding new data to cached data 223 in response to or in anticipation of changes in the behavior of the users of internet 105 and updating cached data 223 in response to changes in source database 241 may of course be implemented in many other ways in the preferred embodiment shown in FIGs. 3 and 4. For example, determining what data should be in cached data 223 could be done in source DBS server 237 instead of in each of the servers 203. Source database 241, like the cached databases 347 in the servers 203(0..n), can maintain statistics information, and a send process 404 in source server 237 can analyze the statistics in substantially the same fashion as described for DSM 323, determine what data should be sent to the servers 203(0..n) for caching in cached data 223, make update queries for that data, and place messages containing the update queries in message queue 414, from which update process 402 can send them to the servers 203.

Updating cached data 223 in response to changes in source database 241 can also be implemented without triggers. The Oracle8 database system includes a redo log buffer 413 in source server memory 415 which is a circular buffer of updates that have been performed on source database 241. The database system maintains the log so that it can redo updates in case of system failure, but the log can also be used to update cached data 223. If there is a table in source database 241 which describes cached data 223, update

process 402 can use the table in conjunction with redo log buffer 413 to determine whether an update in redo log affects cached data 223. If it does, update process 402 can send a copy of the update query to the servers 203 as just described.

5 **Caching servers and source servers that do not involve database systems**

The techniques used to determine what data should be cached in server 203 and to update cached data 223 can also be employed in systems where the data is not queryable. For example, the source data may simply be a collection of documents, identified perhaps by a document number (such as its URL, if the document is an HTML page), and the cached data may be simply a subset of the collection. What cache ~~web-application~~Web application program ~~211-111~~ would receive from HTML component 109 in such a system would simply be the document number for a document; if it is present in the cached data, the caching server would return it from there; otherwise, it would fetch it from the source server. Query log 205 in such a case would be a time-stamped list of the documents that had been requested, together with an indication of whether the document was in the cached data. DSM 213 in such an embodiment would determine as described above for the database whether a document should be included in the cached data, and having made the determination, would obtain it from the source server. As also described above, a send component on the source server could make the same determination and send the document to the caching servers.

For update purposes, the source server would simply maintain a list of the documents that were presently in the caching servers; if one of the documents on the list was updated, updater 243 would send the new version of the document to the caching servers, where DSM 213 would replace any copy of the document in the cache with the new copy. The techniques just described for documents could of course also be used with files and with audio, image, and motion picture data.

Conclusion

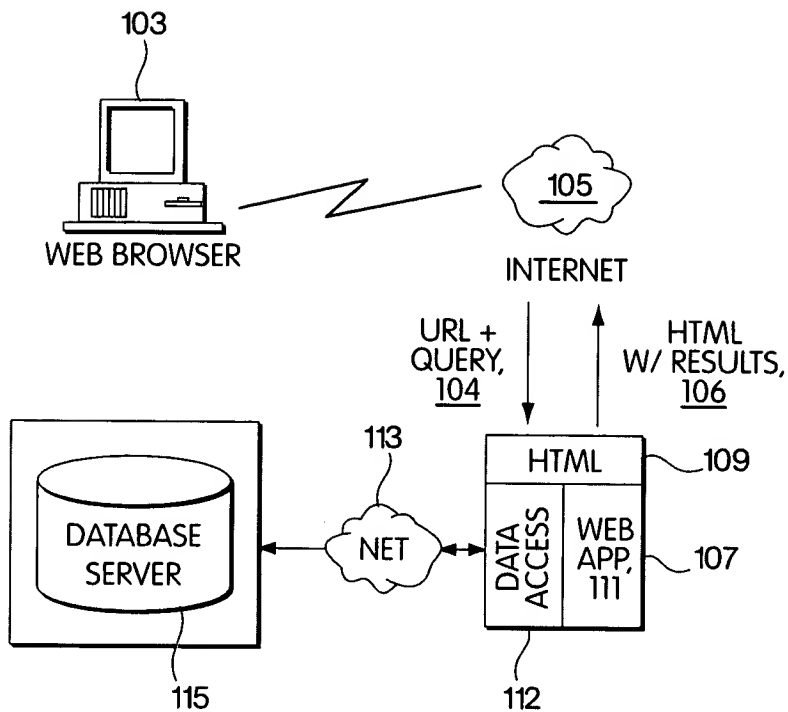
30 The foregoing *Detailed Description* has described a Web server which implements the principles of the inventions set forth herein. The inventions are of course not limited to

Web servers, but may be used in any situation where a cache needs to be kept coherent with the source of the cached data, where there is a need to determine what is going to be cached, where it is desirable to query the cached data, and where it is desired to make the cache transparent to programs running at a higher level. While the inventors have
5 disclosed the best mode presently known to them of implementing their inventions, it will be immediate apparent to those skilled in the arts to which the inventions pertain that there are many other ways of implementing the principles of the inventions.

For all of the foregoing reasons, the *Detailed Description* is to be regarded as being in all
10 respects exemplary and not restrictive, and the breadth of the invention disclosed here in is to be determined not from the *Detailed Description*, but rather from the claims as interpreted with the full breadth permitted by the patent laws.

What is claimed is:

15



101

BROWSER
CLIENT, 103

WEB SERVER, 107 DATABASE SERVER, 115




Fig. 1 (PRIOR ART)




Fig. 4


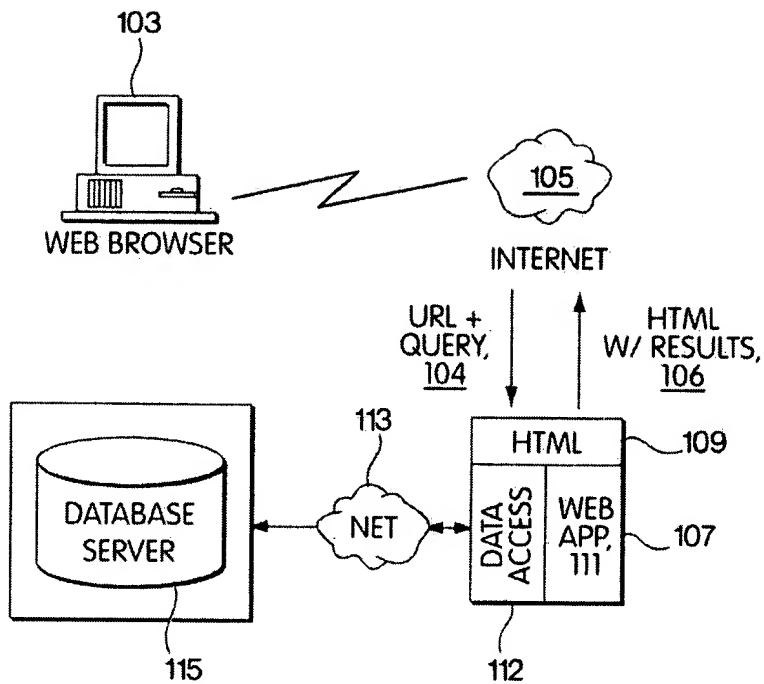


Fig. 5



101

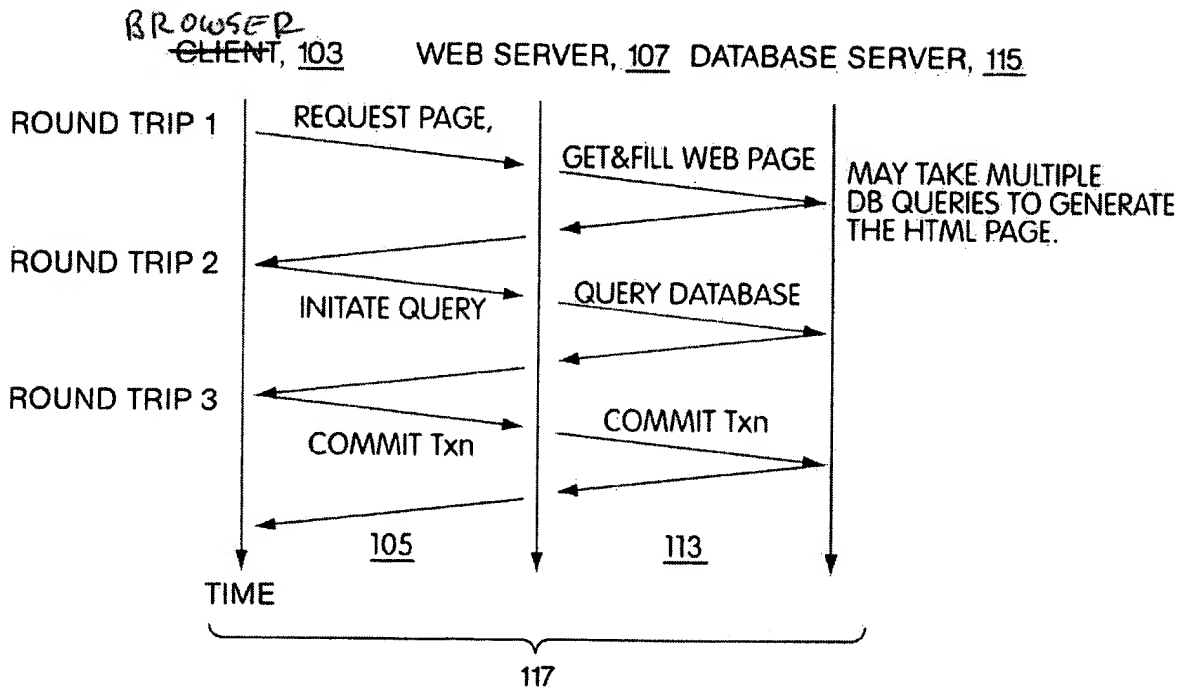


Fig. 1 (PRIOR ART)




Fig. 4




Fig. 5